

Color Non-Uniformity in Projection Based Displays: Analysis and Solutions

Aditi Majumder
Department of Computer Science,
University of North Carolina at Chapel Hill.
majumder@cs.unc.edu

Rick Stevens
Mathematics and Computer Science Division,
Argonne National Laboratories/
Department of Computer Science,
University of Chicago.
stevens@mcs.anl.gov

Abstract

Large-area displays made up of several projectors show significant variation in color. The causes of this color variation are many and complex. The *device-dependent* reasons include the color variations within a single projector's field of view and across different projectors, while the *device independent* reasons include illumination conditions, nature of the display material, shape of the display surface and color bleeding or inter-reflections.

In this paper, we study the device dependent color variation of a multi-projector display and then present a method to solve for it. First, we *identify* the different projector parameters, like position, zoom, axis of alignment, lamp age and projector controls such as brightness, contrast and white balance, that cause such variation. Next, we *study* the change in the luminance and chrominance characteristics of the projection based display with the change in these parameters. From this elaborate analysis we find that the luminance varies significantly within and across projectors while chrominance variation is relatively small, especially across projectors of same brand.

Based on this observation, we present a method to do a per channel per pixel luminance matching. Our method consists of a one-time calibration procedure when a *luminance attenuation map (LAM)* is generated. This LAM is then used to correct any image to achieve photometric uniformity. In the one-time calibration step, we first use a camera to measure the per channel luminance response of a multi-projector display and find the pixel with the most "limited" luminance response. Then, for each projector, we generate a per channel LAM that assigns a weight to every pixel of the projector to scale the luminance response of that pixel to match the most limited response. This LAM is then used to attenuate any image projected by the projector.

This method can be extended to do the image correction in real time on traditional graphics pipeline by using alpha blending and color look-up-tables. To the best of our knowledge, this is the first effort to match luminance across all the pixels of a multi-projector display. Our results show that luminance matching can indeed achieve photometric uniformity.

Keywords: Computer Graphics, Tiled Displays, Projection Based Displays, Color Calibration.

1 Introduction

Large-area, high-resolution multi-projector displays have the potential to change the way we interact with our computing environments. The high resolution and large field of view make them extremely useful for visualizing large scientific models. The compelling sense of presence created by such displays makes them suitable for creating immersive virtual environments for 3D teleconferencing and entertainment purposes. Several such displays exist at Princeton, University of North Carolina at Chapel Hill, University of Minnesota, University of Illinois at Chicago, Stanford, MIT, Fraunhofer Institute (Germany), and U.S. national laboratories such as Lawrence Livermore, Argonne, and Sandia National Laboratories. Recent efforts are directed toward building large displays comprising 40 – 50 projectors (Sandia National Laboratories and National Center for Supercomputing Applications at University of Illinois at Urbana-Champaign).

The color of these displays show significant spatial variation which can be very distracting, thus breaking the illusion of having a single display. The problem of large spatial color variation is unique to these large area projection based displays and has not existed before for other display systems. As has been pointed out in [12,19,20], the causes of this variation are many and complicated. It can be caused by device dependent reasons like *intra-projector color variation* (color variation within a single projector) and *inter projector color variation* (color variation across different projectors) or by other device independent reasons like non-lambertian display surface, curved display surface, inter-reflection and so on. Further, deliberate alignment of projectors in an overlapping fashion to avoid rigid geometric alignment introduces significant color variation across projectors.

Some existing solutions try to reduce the higher brightness in the overlap regions by blending techniques [17] implemented either in software or hardware. But since this does not account for either intra or inter projector variations, the seams between projectors are still visible and one can easily notice the boundaries of the projectors that make up the display, as shown in Figure 1. The solution presented in [12] matches the luminance across multiple projectors but does not account for the variation within a single projector's field of view and hence fails to generate photometrically uniform displays. The comments in the recent works [9, 1, 8, 11, 16, 22] and our experience has led us to believe that this



Figure 1: Left: A display of 5×3 array of 15 projectors where the overlap regions are blended using a physical shadow mask on the light path of the projector. Right: The same display with the overlap region blended by a linear ramp in software. For this, it is necessary to have the knowledge of the exact location of the overlap region. Note that the overlapping regions are distinctly noticeable in both cases.

problem is non-trivial and needs to be analyzed in a structured manner.

In this paper we first *identify* different device dependent parameters that can cause color variation in a large area multi-projector display. Then we *analyze the effects of the changes* in these parameters on this color variation. We also provide insights to the possible reasons for these variations. From this analysis we make the *key observation* that the most significant cause of the spatial variation in color of a multi-projector display is the variation in luminance with the change in projector parameters. Chrominance variation is relatively much smaller, especially for projectors of same brand. Since most multi-projector displays are made of projectors of same brand, this indicates that we may be able to achieve acceptable photometric uniformity by correcting for luminance variation alone.

Finally, we present a method to achieve luminance matching across all pixels of a multi-projector display which results in photometrically uniform displays. We use a camera as measurement device for this purpose. Our method comprises of a one-time calibration step that generates a *per-channel per-projector luminance attenuation map (LAM)* that can then be used to correct any image projected by the projector.

1.1 Main Contributions

Following are the main contributions of the paper.

1. We identify parameters that cause color variations in multi-projector displays and study the effects of changing these parameters on both luminance and chrominance. There has been some work on characterizing specifically the color gamuts of both LCD and DLP projectors [19,20]. Our study in this paper complements this work. Further, the primary goal of our work is to identify, categorize and study the several external and internal parameters that contributes to this variation, within a more general framework.
2. To the best of our knowledge, the camera based algorithm presented here is the first effort that solves for both the intra and inter projector color variations and also the variation introduced by the overlaps. All of these variations are handled by a single algorithm in an automated unified manner which is completely transparent to the user. Some existing solutions match the color properties of different projectors [19,20,12]. However, all these methods rely on taking measurement at only one pixel per projector which indirectly assumes color uniformity within a single projector's field-of-view (intra-projector uniformity). Further, because of this underlying assumption about the intra-projector uniformity, a separate algorithm is needed in the existing methods to handle the overlaps between adjacent projectors. As will be shown by our study and is also acknowledged by other works [19,20,12], this is not true. We present an *automated solution* that would correct for both intra and inter-projector photometric variation within an unified framework.
3. To the best of our knowledge, this is the first effort to achieve photometric uniformity in multi-projector displays that uses a commodity off-the-shelf product like an inexpensive digital camera to measure the spatial luminance variation across the display. Previous work in this direction [12] uses a high-precision expensive radiometer for measurement that takes around 1-20 seconds per reading per pixel. Further, careful positioning of the radiometer is required to read every pixel. This approach becomes impractical when we want to measure the response of potentially each of the millions of pixels on the display.
4. Our method presented in this paper for color correction in projection based displays is easily

scalable, practical and general purpose. Further, this method has the potential to be used in traditional graphics pipelines to achieve this correction in real-time.

This paper is organized in two parts. The first part (Section 2) presents a detailed analysis of the nature of color variation and how different projector parameters are responsible for this variation. From this section, we draw the inference that luminance variation is much more significant in multi-projector displays than chrominance variation. In the second part of the paper (Section 3), we present a method to match the luminance response of every pixel of a multi-projector display using a camera in Section 3.

For the analysis, in Section 2.1 we introduce the basic concepts of color and the desirable color properties of a display. In Section 2.3 we study the *intra* projector color variations. In Section 2.4, we identify the parameters that can affect the color properties of a projector and study the properties of the color variation with changes in these parameters. Next, in section 2.5, we study the *inter* projector color variations. This elaborate analysis indicates that photometric seamlessness may be achieved by correcting for the luminance variation alone.

Thus, in Section 3.1, we present the overview of our algorithm to do such a correction across a multi-projector display. In Section 3.2, we discuss the implementation of each step of the algorithm in details. Then we present the results in Section 3.3. In Section 3.4, we discuss several pertinent issues that would affect the quality of the results achieved by our algorithm. Finally, we conclude with future work in Section 4.

2 Analysis of the Color Variation Properties

In this part of the paper, we present the different experiments we did to study the nature of the color variation across a multi-projector display in details. At the end of this section, we infer some important properties of this color variation from the results of these experiments.

2.1 Background

Color can be specified by three parameters (Y, x, y) . The parameter Y is called the luminance and can be thought of as the amount of achromatic light present in a color. The parameters (x, y) are called chromaticity coordinates and together defines the chrominance.

When two colors $c_1 = (Y_1, x_1, y_1)$ and $c_2 = (Y_2, x_2, y_2)$ are combined additively in proportions p_1 and p_2 (as the primaries are combined in display systems) to give a third color $c_3 = (Y_3, x_3, y_3)$, such that $p_1 + p_2 = 1$, then

$$Y_3 = Y_1 + Y_2; x_3 = p_1x_1 + p_2x_2; y_3 = p_1y_1 + p_2y_2; \quad (1)$$

The three colors (say R,G,B) used to create a display are called primaries, and the independent input paths for these primaries are called channels. The input for each primary has a range from 0.0 to 1.0. The colors projected by the display for the input of 1.0 at each channel (and 0.0 in other two channels) be $(Y_R, x_R, y_R), (Y_G, x_G, y_G)$ and (Y_B, x_B, y_B) respectively. The triangle formed by the chromaticity coordinates of these primaries is called the *color gamut* of the display. Readers are referred to [6] for additional in-depth treatise on colorimetry.

Ideally, it is desirable to have a display where given the properties of the primaries one can predict, using simple formulae, the properties of any color produced by the combination of the primaries. This becomes easy if the display satisfies the following properties.

1. *Channel Independence* : This assumes that the light projected from one channel is independent of the other two. Hence, this assumes that there is no leakage light from other channels to interfere with the light projected from a channel.
2. *Channel Constancy* : This assumes that for each channel, with changing inputs only luminance changes while the chromaticity coordinates remain constant. This means that for input $0.0 \leq r \leq 1.0$, $x_r = x_R$ and $y_r = y_R$ and Y_r changes, where (x_r, y_r) is the chromaticity coordinate of r and Y_r is the luminance of r .
3. *Spatial Homogeneity* : The response of all the pixels of the display device is identical for any given input.
4. *Temporal Stability* : The response for any input at any pixel of the device does not change with time.

The property of optical superposition states that light falling at the same physical location from different sources adds up. The properties of channel constancy, channel independence, optical superposition along with the assumption that with an input of $(0, 0, 0)$ the display outputs zero light indicates

that the color projected at a pixel is a linear combination of the color projected by the maximum values of the red, green and blue channels alone when the values of the other two channels are set to zero. Hence, for any input $c = (r, g, b)$, $0.0 \leq r, g, b \leq 1.0$, the luminance Y_c is given by $Y_r + Y_g + Y_b$ and the chromaticity coordinate is given by the barycentric coordinate

$$x_c = rx_R + gx_G + bx_B; y_c = ry_R + gy_G + by_B; \quad (2)$$

This is referred to as the *linear combination property*.

Given the linear combination, the spatial homogeneity, and the temporal stability property, it is easy to predict the color properties of any pixel of the display surface for any input if the response of the primaries at any one pixel of the display is known. Most traditional display devices like CRT monitors satisfy these properties to a reasonable accuracy or the deviation from this ideal behavior is simple enough to be modelled by simple linear mathematical functions [4]. However, as we will see in the following sections, a projector is not such an ideal device.

However, before we delve deep into the study of the property of the color variation, we would like to talk about the measurement process and instruments and other extraneous factors that may have an effect on the measurement process.

2.2 Measurement

As test samples for our measurement process, we used multiple Sharp, Nec, Nview, Proxima and Epson projectors including both LCD and DLP projectors. We have taken measurements for both front and back projection systems.

2.2.1 Measuring Devices

The goal of the process of measurement is to accurately find the luminance and the chrominance properties at different points of the tiled display. There are two options for the optical sensors that one might use for this purpose.

1. **Spectroradiometer :** This is an expensive precision-instrument giving high accuracy measurements. But, it can measure only one point at a time. It is very slow taking about 1 – 20 seconds for each measurement. Further, it is difficult to position is accurately to read every pixel separately. This makes it unsuitable for acquiring high resolution data in a reasonable amount of

time with a high geometric accuracy. The advantage, however, lies in its high photometric accuracy. It can measure any color in the visible spectrum by breaking up the incoming light into several spectral bands and then calculating its spectral distribution. Hence, it can measure any color that the projector is capable of projecting, enabling us to work in a laboratory-calibrated device-independent color space.

2. **Camera :** The camera, on the other hand, is relatively inexpensive and is capable of measuring a large number of points in a single snap. Thus it is suitable for acquiring high resolution data in a very short time. There are many existing algorithms to find the correspondences between the camera coordinates and the projector coordinates we are measuring, which assures geometric accuracy of the data acquired. But, there are some limitations on the photometric accuracy that can be achieved by using a camera. First, we cannot measure all the colors projected by the projector, if the color gamuts of the camera does not contain the gamut of the projector, thus restricting us to work in a device dependent color space. Second, this is not a linear device and we need to reconstruct the camera non-linearity to measure the projector response reliably with it. Third, the camera may have spatial inhomogeneity itself which needs to be factored out from the measurements. Finally, the exposure of the camera needs to be adjusted so that it is not produce over or under-saturated measurements.

With these options in hand, we used both types of sensors, but for different purposes. For point measurements, we used a precision spectroradiometer (Photo Research PR 715). But, for finding the spatial color/luminance variation across a projector, we need to measure the color at potentially every pixel of a projector. Hence, we use a high resolution digital camera for these readings. To assure reliable measurements, we reduce the photometric inaccuracies introduced by the camera by a reasonable amount by the following methods.

2.2.2 Using Camera as Reliable Measuring Device

In our approach we take care of the four most important issues related to using camera to measure projected colors. They are non-linearity of the camera response, flat-field response, dynamic range, and projector to camera pixel correspondence. Since we cannot change the internal color filters of the camera, there is no way to enlarge the camera color gamut.

The non-linearity of the camera is recovered using the algorithm presented in [3]. From this we generate a color look-up-table that would linearize the camera response. Every picture from the camera is linearized using this color look-up-table.

It is important that the camera does not introduce additional luminance variation than is already present on the display wall. Hence, the camera must produce flat fields when it is measuring a flat color. It is mentioned in [3] that most cameras satisfy this property at lower aperture settings, especially below $F8$. Our camera had a standard deviation of 2 – 3% for flat field images. These flat field images were generated by taking pictures of nearly-diffused planar surfaces illuminated by a studio light with a diffusion filter mounted on it.

It is important for the brightness of the projector to be well within the dynamic range of the camera. We verify this by simple under or over saturation tests of the camera images. The exposure should be taken into account by appropriate scaling factors while generating the luminance surface [3].

Finally, for the camera to projector correspondence, we use a geometric calibration method [7] to find the mapping from projector pixel to its corresponding camera pixel. This assures the geometric accuracy of the measurements from the camera.

It will be evident from the following sections that we use the camera for measuring the luminance or chrominance at a high spatial resolution. To extract the luminance or chrominance from camera color data, we use standard linear RGB to YUV transformation.

2.2.3 Screen Material and View Dependency

It is ideal to use a diffused lambertian screen so that the screen properties do not amplify the color variation. For the front projection screen experiments we use such a screen. But the Jenmar screen we use for our back projection system is not exactly lambertian. This results in the measuring devices being sensitive to viewing angles. In case of the spectroradiometer, which is used to take measurement at single point at a time, we orient it perpendicular to the point that is being measured. But for the camera, the view dependency cannot be eliminated. However, as it will be evident from the later sections, we use the camera to study the *shape* of the luminance variations. And, we do not see any inconsistencies in the shapes of the measured curves and surfaces due to this view dependency.

2.2.4 Ambient Light

Both the spectroradiometer and the camera can see all the ambient light. Hence we try to reduce stray light as much as possible by taking the readings in a dark room turning off all lights. When taking the measurement of a single projector, we turn off all the adjacent projectors. For our spectroradiometer measurements of the front projection systems, we used black material to cover up the white walls of the room to avoid inter-reflected light.

In the next few sections we study the intra and inter-projector color properties from the measurements taken using a spectroradiometer or a camera.

2.3 Intra Projector Variations




Figure 2: Left: Luminance Response of the three channels; Right: Chromaticity x for the three channels. The shape of the curves for chromaticity y are similar

First, we study the intra projector variations. In the process we will show that the projectors do not follow the desirable properties mentioned in Section 2.1. A few of these results are also confirmed in [13,19,20].

One important consequence of a display to satisfy *channel independence* and *channel constancy* property is that the response for black (input of $(0, 0, 0)$) should have zero light. However in projectors, because of leakage light, some light is projected even for black. This is called the *black offset*. From the various measurements we had from the spectroradiometer, we found that this can be up to 2%




Figure 3: Left: Gamut Contour as the input changes from 0 to 255 at intervals of 32; Right: Color gamut at four different spatial locations




Figure 4: Left: Luminance Response of the red channel plotted against input at four different spatial locations; Right: Luminance Variation of different inputs of the red channel plotted against spatial location. The responses are similar for other channels.

of the maximum luminance projected per channel. Hence the chromaticity for any channel at zero is the chromaticity of this achromatic black. As the inputs increase, the chromaticity reaches a constant value as it should for a device following channel constancy. This is demonstrated in Figure 2. The contours shown in Figure 3 shows how the gamut starts out as a single point for 0 in all the three channels and then attains the final red triangle at the highest input value. However, this black offset can be modelled as a linear offset term to be subtracted from the response of all inputs as has been shown in [20].

If the black offset is accounted for by the linear offset term, almost all projectors exhibit linear combination property in their responses to the input except some DLP projectors, which do not exhibit the linear combination property for the grays. We found that in these projectors, a clear filter is used while projecting the grays, instead of adding up the contributions from the red, green and blue channels. Hence, luminance of the grays are much higher than the sum of the luminances of the constituting red, green and blue. The effect has been modelled in [19] by an additive gamut with an extrusion at the white point.

Projectors are not spatially homogeneous either. Accurate luminance and chrominance readings were taken at equally spaced five locations on the projector diagonal using the spectroradiometer. We named these locations from 1 to 5 starting at the top left corner position. The luminance reaches a peak at the center (location 3) as seen in Figure 4. The luminance falls off at the fringes by a factor which may be as high as 80% of the peak luminance at the center for the rear projection systems, and about 50% for front projection system. We believe that the non-lambertian nature of the display makes this luminance fall-off view-dependent in rear-projection systems. And our measuring devices are sensitive to such view dependencies. However, this considerable fall-off in luminance indicates that having wide overlaps between projectors in a multi-projector display can help us to get a better overall dynamic range.

Further, the important thing to note here is, only the luminance differs while the color gamut remains almost identical as shown in Figure 3. The gamut is measured from the chromaticity coordinates of the primaries at their highest intensities. However, this does not assure that the gamut at different spatial locations do not differ at lower input values, which in turn, may indicate that the chrominance varies within a single projector’s field of view for some inputs. We confirmed this is *not* the case by finding the per channel chromaticity response for every input at different spatial locations using the

spectroradiometer which were nearly identical.

Given these observations from the spectroradiometer measurements, we used a camera to measure the intra-projector spatial luminance variation at a much higher resolution. First, in most projectors the peak does not appear at the center but is skewed a little around center. The basic commonality is that there is a peak luminance response somewhere near the center and then the luminance falls off radially towards the fringes. Second, attenuation is not symmetric. Finally, the location of the center and fall off pattern can vary from projector to projector. Figure 17 shows the luminance variation across a single projector’s field of view.

These observations can be explained easily. The chrominance depends on the physical red, green and blue filters of the projectors. Since the filters do not change, the chrominance should not change within a single projector’s field of view. On the other hand, the intensity of light from a point light source falls off as the distance from the center increases which is the cause of the luminance variation. In addition to this, the non-lambertian nature of the display makes the luminance variation more acute. The asymmetry in the fall-off pattern gets pronounced with off-axis projection, as we will see in the following sections. This indicates that the orientation of the projector is responsible for this asymmetry.

Finally, projectors are not temporally stable. The lamp in the projector ages with time which changes the color properties of the projector. Figure 5 show the results. Here also, we see a significant difference in luminance even within a short amount of time while the chrominance remains almost same. However, we have found the color characteristics also drifts a little, but only after extensive use of about 800 – 900 hours. This indicates that the lamp initially changes in luminance, but over longer periods of time there is also a shift in its color characteristics.

2.4 Projector Parameters that Change Color Properties

Now that we understand the intra projector variations in luminance and chrominance, in this section, we will identify the different projector parameters that can change the color properties of a projector, study the effects of varying these projector parameters on the color properties of a large area multi-projector display and provide insights for the possible reasons behind such effects.




Figure 5: Left: Luminance Response of the green channel at four different bulb ages; Right: Chrominance Response of the green channel with four different bulb ages. Note that all the curves are coincident indicating that chrominance changes negligibly with short periods of time.

2.4.1 Position

Position defines the location and the orientation of the projector with respect to the display wall. This has two components, the distance from the wall along the axis of projection and the alignment of the axis of projection with the planar display surface. We did two sets of experiments. In one we kept the orientation of the projector constant and moved it at different distances from the wall along its axis of projection. In the other, the axis of projection was changed while keeping its distance from the wall constant.

1. **Distance form the Wall:** We moved the projector at different positions along its axis of projection and measured its luminance and chrominance properties. The results are shown in Figure 6. As you notice, the luminance changes while the chrominance remains constant. Further, the shape of the spatial variance of the luminance also remains the same as shown in Figure 4 and 17. By moving the projector away from the wall, the projection area increases. Hence the amount of light falling per unit area can change, but the nature of the fall off should not change. This is reflected in the observation that the average luminance decreases while the nature of the luminance fall-off remains unaffected with the change in the distance from the wall.




Figure 6: Left: Luminance Response of the green channel as the distance from the wall is varied along the axis of projection; Right: Chrominance Response of the green channel as the distance from the wall is varied along the axis of projection. Note that the curves are coincident indicating the chrominance does not change with the change in the distance of the projector from the wall




Figure 7: Left: Luminance Response of the red channel with varying axis of projection; Right: Luminance response of different inputs in the green channel plotted against the spatial location along the projector diagonal.

2. **Off-Axis or Orthogonal Projection:** In this set of experiments, we kept the projector at the same distance from the wall while we rotated it about x , y , and z direction to have a off-axis projection and study how this affects the luminance and chrominance properties. Figure 7 shows the result for four orientations between orthogonal to 60 degrees angled axis of projection. In this case also we found that the chrominance remains constant while the luminance response changes. Hence we show only the luminance response here. The nature of the spatial variation is no longer symmetric as shown for an orthogonal position in Figure 4. Near the longer boundary of the key-stoned projection which is physically farther away from the projector, there is a higher drop in luminance as shown in Figure 7. As the orientation becomes more oblique, the luminance attenuation at the projector boundary further away from the screen increases, resulting in asymmetric luminance attenuation. This is due to two reasons. First, there is larger attenuation due to larger distances at the farther boundaries, and second, the light from each pixel gets distributed over a larger area. The results are similar in the both the horizontal and vertical direction.

In both the above cases, since moving the projector around does not change the internal filters of the projector, the chrominance remains constant as is expected.

2.4.2 Controls

The projectors offer us various controls like zoom, brightness, contrast and white balance. Knowing how these controls affect the luminance and the chrominance properties of the projector can help us to decide the desirable settings for the projector controls to help reduce variation within and across projectors and thus, to avail of the best possible dynamic range and color resolution offered by the device.

1. **Zoom:** We tested the projector for four different zoom settings. We found that both luminance and chrominance remains constant with the change in zoom settings of the projector. With the change in zoom, the amount of light for each pixel gets distributed over a different area. For a focussed projector, it is distributed over a small area, while for a unfocussed projector it is distributed over a larger area. However, the total area of projection remains the same and the total amount of light falling in that area remains same. Hence the light per unit area remains

unchanged, while the the percentage of light that each unit area receives from the different pixels changes.

- Brightness:** Luminance and chrominance response was measured by putting the brightness control in 5 different positions. It is mentioned in [14], usually the brightness control is used to change the black offset. However, we found that in projectors, this control affected both the gain and black offset of the luminance response of *all the three channels similarly and simultaneously*. As the brightness is increased, both the black offset and the gain of the luminance increased. However, if the brightness is too low, the luminance response gets clipped at the lower luminance range. In these settings, since the luminance remains at the same level for many lower inputs, the chromaticity coordinates also remains constant at gray for these inputs. For some projectors, at very high brightness settings, we observed some non-monotonicity in the luminance range for higher intensity input range. As a consequence, the chromaticity coordinates also show some instability at the higher brightness settings. Figure 8 and Figure 10 illustrates these effects. It is ideal to have the brightness control set so that there is no clipping in the lower input range or non-monotonicity at higher input ranges. For example, in these particular illustrations, the ideal setting is between 0.5 and 0.75.




Figure 8: Left: Luminance Response of the green channel with varying brightness settings; Right: Luminance Response of the green channel with varying brightness settings zoomed near the lower luminance region to show the change in the black offset.

3. **Contrast:** We performed similar experiments for the contrast control. This also affects *all the three channels similarly and simultaneously*. As mentioned in [14], usually the contrast control is used to change the gain of the luminance curve. We found the same with the projectors as illustrated in Figure 9. While the gain is affected, the black offset remains the same. As the gain increases, the luminance difference became significant enough at lower input ranges to push the chromaticity away from the gray chromaticity values towards the chromaticity coordinates of the respective primaries, as illustrated in Figure 10. However, the luminance response starts to show severe non-monotonicity at higher contrast settings, thus reducing the input range for which the luminance shows monotonic behavior, as illustrated in Figure 9. This shows that the gain may not be increased beyond a limit after which clipping starts to happen. Hence it is important to keep the contrast setting in the monotonic range to maximally use the available color resolution.




Figure 9: Left: Luminance Response of the green channel with varying contrast settings; Right: Luminance Response of the green channel with varying contrast settings zoomed near the lower luminance region to show that there is no change in the black offset.

4. **White Balance:** The white balance usually has a brightness and contrast control for each of the three channels separately. When we studied the color response by putting these in five different settings, we found the luminance and the chrominance response changes exactly the same way as for the independent brightness and contrast controls, but the changes affect only one channel at a time instead of affecting all of them similarly. Hence this can be used to control the proportion




Figure 10: Left: Chrominance Response of the green channel with varying brightness settings; Right: Chrominance Response of the green channel with varying contrast settings.

of the contribution from each channel to a color which in turn changes the white balance. This is illustrated in Figure 12 and 11.

2.5 Inter Projector Color Variations

In the previous sections we studied the intra-projector variations and how projector parameters affect them. The projection based systems are made of multiple projectors and hence it is important to know how these properties vary from one projector to another.

Figure 13 shows the luminance and color gamut response for the maximum intensity of a single channel for different projectors of *same brand* having exactly the same values for all the parameters defined in Section 2.4. There is nearly 66% variation in the luminance, while the variation in color gamut is relatively much smaller. In Figure 14, we show a high resolution chrominance response of a display wall made of four overlapping projectors of same brand to emphasize that the chrominance difference is indeed very small. This difference in chrominance is within the just noticeable difference threshold designed in [5]. Projectors of same brand and make usually use same brand bulb (which have similar white points) and have similar filters. Hence, we do not expect to see much variation in color gamut.

While the color gamut of the projectors of same brand is almost identical, the color gamut differs




Figure 11: Left: Chrominance Response of the green channel with varying green brightness settings for white balance; Right: Chrominance Response of the red channel with varying red contrast settings for white balancing.




Figure 12: Left: Luminance Response of the green channel with varying red brightness settings in white balance does not change; Right: Luminance Response of the red channel with varying red brightness settings in white balance varies.




Figure 13: Left: Peak luminance of green channel for fifteen different projectors of the same brand. Right: Color gamut of 5 different projectors of the same brand. Notice the large variation in luminance and small variation in chrominance.




Figure 14: Left: Chrominance response of a display wall made of four overlapping projectors of same brand. Right: Color gamut of projectors of different brand.

<i>Projector Brand</i>	<i>Red</i>		<i>Green</i>		<i>Blue</i>	
	<i>x</i>	<i>y</i>	<i>x</i>	<i>y</i>	<i>x</i>	<i>y</i>
Sharp XG-E3000U	0.62	0.32	0.33	0.62	0.14	0.07
NEC MT-1035	0.55	0.31	0.35	0.57	0.15	0.09
nView D700Z	0.54	0.34	0.28	0.58	0.16	0.07
Epson 715c	0.64	0.35	0.30	0.67	0.15	0.05
Proxima DX1	0.62	0.37	0.33	0.55	0.15	0.07
<i>Max Distance</i>	0.085		0.086		0.028	

Table 1: Chromaticity Coordinates of the primaries of different brands of projectors a little across projectors of different brands as shown in Table 2.4.2. This is also illustrated in Figure 14. However, this is extremely small when compared with the luminance variation.

2.6 Inference

The observations from experiments and analysis of section 2.4 and 2.5 can be summarized as

- Within a single projector’s field of view, only luminance varies while chrominance remains almost constant.
- Across different projectors of same brand, chrominance variation is negligible while luminance variation is significant.
- Chrominance varies a little across projectors of different brand, but this variation is very small when compared to the variation in luminance.
- With the change in various projector parameters like brightness, contrast, zoom, distance and orientation, only luminance changes while chrominance remains constant.

Thus, the key observation is that for every channel, *luminance undergoes significant variation while chrominance and the color gamut does not vary much*. Further, in almost all cases, multi-projector display walls do comprise of projectors of the same brand. In addition to this, from a perceptual standpoint, the humans are more sensitive to luminance changes than to chrominance changes [21].

All these indicate that correcting for the luminance across a multi-projector display may be sufficient to achieve photometric uniformity.

3 Luminance Matching Algorithm

Aided with the key observation from the preceding section, that luminance matching can achieve photometric uniformity, in this section, we present a method that matches the luminance at every pixel of a multi-projector display with each other.

3.1 Algorithm Overview

In this section, the algorithm is described for a single channel. All three channels are treated similarly and independently.

The method comprises two steps. The first step is a one-time calibration step where a per projector luminance attenuation map is generated. In the second step, this LAM is used to correct any image content.

3.1.1 Calibration Step

The calibration step consists of three stages.

1. **Measuring the Luminance Response:** The *luminance response* of any pixel is defined as the variation of luminance with input at that pixel. We measure the luminance response of every pixel of the display with a camera.
2. **Finding the Common Achievable Response:** We find the common response that every pixel of the display is capable to achieving. The goal is to achieve this *common achievable response* at every pixel.
3. **Generating the Luminance Attenuation Map:** We find a luminance attenuation function that transforms the measured luminance response at every pixel to the common achievable response.

If we assume a linear response for the projectors, then each of these three stages gets simplified. By linear response we mean that the luminance of black is zero, the maximum luminance occurs for

the maximum input, and luminance response for every other input is a linear interpolation between these two values. First, the luminance measurement stage is simplified with this assumption because instead of measuring the luminance response of every input, we can now measure the luminance of only the maximum input. Second, the common achievable response can now be defined as the linear response with minimum luminance range. Third, the luminance attenuation function is just a scaling function that is encoded in the luminance attenuation map. Hence, we assume that every display pixel has a linear luminance response. In Section 3.2 we show how we satisfy this assumption in the actual implementation.

1. Measuring the Luminance Response

Let us assume that the display D of resolution $W_d \times H_d$ is made up of n projectors each of resolution $W_p \times H_p$. Let us refer to the projectors as $P_i, 0 \leq i < n$. We use a static camera C of resolution $W_c \times H_c$ to measure the luminance of D . Let us denote the luminance response for the *maximum input* of the channel at a display location (x_d, y_d) as $L_d(x_d, y_d)$. The light at (x_d, y_d) can come from one or more projectors. If it comes from more than one projector, then (x_d, y_d) is in the region of the display where multiple projectors overlap. We want to find $L_d(x_d, y_d)$ for all pixels (x_d, y_d) .

Geometric Calibration: First, we perform a geometric calibration that defines the geometric relationships between the projector pixels (x_{P_i}, y_{P_i}) , camera pixels (x_c, y_c) and the display pixels (x_d, y_d) . This geometric calibration uses the static camera to take pictures of some known static patterns projected on the display. By processing these pictures, the geometric calibration procedure defines two warps : $T_{P_i \rightarrow C}(x_{P_i}, y_{P_i})$, which maps a pixel (x_{P_i}, y_{P_i}) of projector P_i to the camera pixel (x_c, y_c) , and $T_{C \rightarrow D}(x_c, y_c)$, which maps a camera pixel (x_c, y_c) to a display pixel (x_d, y_d) . The concatenation of these two warps defines $T_{P_i \rightarrow D}(x_{P_i}, y_{P_i})$, which maps a projector pixel (x_{P_i}, y_{P_i}) directly to display pixel (x_d, y_d) . These three warps give us the geometric information we need to find $L_d(x_d, y_d)$.

Data Capture for Luminance Correction: Keeping the camera in the same position, we take the image of each projector P_i projecting the maximum input for the channel. From these images




Figure 15: The common achievable response with four sample pixel response. The response with the least range is the common achievable response.

we extract the luminance image, denoted by I_i , for each projector P_i in the camera coordinate space.

Generation of the Luminance Surface: Next we generate the luminance surface $L_{P_i}(x_{P_i}, y_{P_i})$ for every projector P_i . For this, we first transform every projector pixel (x_{P_i}, y_{P_i}) by $T_{P_i \rightarrow C}$ into the camera coordinate space and read the luminance at that transformed pixel from I_i . Hence

$$L_{P_i}(x_{P_i}, y_{P_i}) = I_i(T_{P_i \rightarrow C}(x_{P_i}, y_{P_i})) \quad (3)$$

Once we have the luminance surface L_{P_i} for every projector P_i , we find the contribution of every projector at (x_d, y_d) by the inverse warp of $T_{P_i \rightarrow D}$ denoted by $T_{D \rightarrow P_i}(x_d, y_d)$ and add them up.

$$L_d(x_d, y_d) = \sum_{i=1}^n L_{P_i}(T_{D \rightarrow P_i}(x_d, y_d)) \quad (4)$$

2. Finding the Common Achievable Response

The common achievable response is defined as a linear response for which the luminance response for the maximum input is minimum of all $L_d(x_d, y_d)$ and this minimum luminance is denoted by L_{min} . Conceptually, this is equivalent to finding a common response that every pixel is capable of achieving. Figure 15 illustrates this.

3. Generating the Luminance Attenuation Map

The LAM, denoted by $A_d(x_d, y_d)$, is first generated in the display coordinate space and is given

by

$$A_d(x_d, y_d) = \frac{L_{min}}{L_d(x_d, y_d)} \quad (5)$$

Thus A_d signifies the pixelwise scale factor (less than 1.0) by which L_d should be scaled down to achieve luminance matching.

The next step is to generate the per projector luminance attenuation maps $A_{P_i}(x_{P_i}, y_{P_i})$ from A_d . Since we know the warp $T_{P_i \rightarrow D}$, this is achieved by

$$A_{P_i}(x_{P_i}, y_{P_i}) = A_d(T_{P_i \rightarrow D}(x_{P_i}, y_{P_i})) \quad (6)$$

3.1.2 Image Correction Step

Once this per projector LAM is generated, it is used to attenuate any image. When an image $M(x_d, y_d)$ of resolution $W_d \times H_d$ is projected on the display wall, the warp $T_{P_i \rightarrow D}$ is used to generate $M_{P_i}(x_{P_i}, y_{P_i})$ which is the part of M that projector P_i should project.

$$M_{P_i}(x_{P_i}, y_{P_i}) = M(T_{P_i \rightarrow D}(x_{P_i}, y_{P_i})) \quad (7)$$

Finally, M_{P_i} is multiplied by A_{P_i} to create the final image for projector P_i , denoted by F_{P_i} .

$$F_{P_i}(x_{P_i}, y_{P_i}) = M_{P_i}(x_{P_i}, y_{P_i}) \times A_{P_i}(x_{P_i}, y_{P_i}) \quad (8)$$

3.2 Implementation

In this section we will describe how our method is implemented. The implementation is done on two wall configurations. The first one is a wall of resolution 1200×800 made up of 2×2 array of four projectors. Later we extended this to a wall of resolution 4500×2000 made of 5×3 array of fifteen projectors.

3.2.1 Luminance Response Measurement

This section focuses on the luminance response measurement.

1. Geometric Calibration

We need an accurate geometric calibration algorithm for our photometric calibration. Several geometric calibration algorithms have been designed in the past [16, 15, 22]. Any geometric calibration algorithm that can define accurately the two warps, $T_{P_i \rightarrow C}$ and $T_{C \rightarrow D}$, can be used for our method. For our implementation, we use two *cubic nonlinear* warps to define $T_{P_i \rightarrow C}$ and $T_{C \rightarrow D}$. These non-linear warps include the radial distortion correction for both the camera and the projectors and can be implemented in real time on traditional graphics pipeline by using texture mapping. The details of our algorithm are available in [7].

2. Data Capture for Luminance Correction




Figure 16: Top : The four pictures taken for green channel to do the luminance attenuation for a display made of 2×2 array of 4 projectors. Bottom : The four pictures taken for green channel to do the luminance attenuation for a display made of 5×3 array of 15 projector.

As mentioned in the preceding section, we need to capture images for every projector P_i when it is projecting the maximum input for each channel. During this time we turn off all the projectors that overlap with P_i to capture the luminance contribution solely from P_i accurately. To capture the data for all projectors in the display, we need to take a total of four pictures per channel. In each picture alternate projectors are turned on so that none of them overlap with each other. The pictures taken for the two different wall configurations are shown in Figure 1 of the color page.

In the preceding section, we assumed linear devices for our algorithm. To satisfy this assumption, we find the camera's nonlinear response and linearize it using a color look-up-table. For our implementation we use a Fujifilm MX-2900 camera. We use the method presented in [3] to




Figure 17: Left : The luminance surface generated for one projector. Right : The same luminance surface after edge attenuation.




Figure 18: Left : The luminance surface generated for 2×2 array of four projectors. Right : The luminance surface generated for 5×3 array of 15 projectors.

reconstruct its nonlinear response. This method generates a per channel color look-up-table that linearizes the per channel luminance response of the camera. Every image captured by the camera is linearized using this look-up-table.

3. Generating the Luminance Surface

Generating the luminance response surface for the display requires several steps.

Generating the Luminance Surface in Camera Coordinate Space: First, we find the luminance surface in the camera coordinate space corresponding to linearized images generated in the preceding section. For this we use the standard linear transformation usually used to convert RGB colors to YUV space given by

$$Y = 0.299R + 0.587G + 0.114B \quad (9)$$

Generating the Per Projector Luminance Surface: In this step, we generate L_{P_i} for each projector P_i . For every pixel of the projector we find the corresponding camera coordinate using $T_{P_i \rightarrow C}$ and then interpolate bilinearly the corresponding luminance from the luminance of the four nearest neighbors in the camera coordinate space. An example of the luminance surface thus generated for a projector is shown in Figure 17.

Edge Attenuation: In most projection based displays, adjacent projectors are overlapped to avoid rigid geometric alignment. However, the luminance in the overlap region is much higher than the luminance in the non-overlapped region and this spatial transition is very sharp. Theoretically, to reconstruct this edge between the overlapped and non-overlapped regions we would need a camera resolution at least twice the display resolution. Given the resolution of today’s display walls, this is a severe restriction.

Instead, we smooth out this sharp transition by attenuating a few pixels at the edge of each projector. This increases the error tolerance to inaccuracies in reconstruction of the luminance surface in regions of sharp transition. We do this attenuation in software. After generating the luminance image for each projector, we attenuate the 40 – 50 pixels at the edge of the projector using a linear function. (The width of this attenuation can be changed as long as it is less

than the width of the overlap region. Similarly, a different function can be used e.g. a cosine ramp.) Figure 17 shows the luminance after such an edge attenuation. Note that we do not need information about the exact location of the overlap regions for this purpose but just an approximate idea about the width of the overlap so that the attenuation width is less than the width of the overlap. Further, this approach allows us to process the luminance of each projector independently, without explicitly considering geometric correspondences across the projectors.

Adding Them Up: Now, we have got the luminance image for each projector. The next step is to add them all up in the display coordinate space to generate L_d . For every projector pixel, we use $T_{P_i \rightarrow D}$ to find the corresponding display coordinate and then add the contribution of the luminance to the nearest four display pixels in a bilinear fashion. The generated luminance surface for the 2×2 array of four projectors and the 5×3 array of fifteen projectors is shown in Figure 18.

3.2.2 Luminance Attenuation Map Generation




Figure 19: LAM for a display made of the 5×3 array of 15 projectors.

We define the common achievable response as the minimum of L_d designated by L_{min} . Then we generate the luminance attenuation map A_d , in the display coordinate space by dividing L_{min} by L_d . This is shown in Figure 19. Notice how the LAM is dimmer to compensate for the brighter regions of

the luminance surface in the overlap regions and near the center of each projector. Further, because of the large luminance fall-off at the edges of the boundary projectors where there is no overlap, the reduction in dynamic range can be drastic leading to unacceptable picture quality. Hence, we ignore about 200 pixels in the boundary of the display coordinate system while generating the LAM.

To generate the per projector attenuation map A_{P_i} , for every projector pixel we use $T_{P_i \rightarrow D}$ to convert it to display coordinate space and then interpolate bilinearly the value of A_d from the nearest four neighbors.

Finally, we put in the edge attenuation in the luminance attenuation map for each projector by attenuating the same number of edge pixels in the same way as was done while generating the luminance image in the preceeding section. Figure 20 shows an example LAM for one projector. The fifteen projector wall had larger luminance variation, with some of the projectors having very low luminance response. Hence the attenuation in the fifteen-projector display is higher than that in the four-projector display.




Figure 20: Left : LAM for a single projector in the four projector display. Right : LAM for a single projector in the 15 projector display.

3.2.3 Image Correction

The image correction is done in two steps.

Image Attenuation: The LAM is multiplied with the image to be rendered. This can be extended to an interactive application using traditional graphics hardware, where the LAM can be used as an

alpha mask that is blended with the rendered image.

Linearization of Projectors : Since we have assumed linear response for the projectors, we have to linearize the projectors. This is done by a look-up-table. These per projector look-up-tables are pregenerated. It is shown in [13] that the projector non-linearity response does not vary spatially. Hence, we use a photometer to measure the per channel nonlinear luminance response only at the center of every projector. Then we find a look-up-table that would linearize this luminance response and use this for all pixels of the projector.

3.3 Results

In this section we present and discuss our results in the four and 15 projector display walls. Figure 21 and 22 shows the results on the four projector wall. These images are taken by a digital camera using the same exposure so that they can be compared. The worst test patterns for this algorithm are images with flat test colors. Figure 22 shows our algorithm on such images. A faint vertical line that can be seen in the images is not the projector boundaries but is the physical crack between the vertical planks that make our display screen.

Figure 23 shows results of the 15 projector display. You will see two types of artifacts in these results. Some contours are visible and some of the projector edges are faintly visible. These artifacts are due to insufficient sampling or limited camera dynamic range and will be explained in details in the next section. The bright spots you see in the center are due to light leaking through the cracks between the planks making up the display. Due to larger variation in luminance, the attenuation is larger for the 15 projector display. Hence, the images of the corrected display are taken at a higher exposure than the images of the uncorrected display.

The LAM can be implemented using the conventional graphics pipeline in real time by alpha blending. However, for the final linearization in the image correction step, we need a LUT. Usually all off-the-shelf projectors have in-built hardware LUT which would be ideal for this purpose since this would not incur any extra computation overhead. However, most commercial projectors do not give the user complete access to this hardware LUT. Hence we had to implement this using the software LUT in OpenGL. Unfortunately, this becomes the bottleneck stage in terms of achieving interactive speeds. We can render a movie using OpenGL at 15 frames per second just with the alpha mask. Note



Figure 21: The left column shows the image before correction and the right column shows the image after luminance matching.




Figure 22: The left column shows the image before correction and the right column shows the image after luminance matching.



Figure 23: The left column shows the image before correction and the right column shows the image after luminance matching.

that this speed is limited by the time required to load the movie and not to render it. However, it does not look right without the final linearization. But, if we use the OpenGL LUT for this purpose, it takes 2 – 3 seconds per frame on nVidia GeForce2 cards. Currently we are trying to find some projectors that would give us complete access to their hardware LUTs so that we can implement an interactive version of this algorithm.

3.4 Issues

As a result of our work, we have identified several issues that we now comment on.

Accuracy of Geometric Calibration: Our geometric calibration algorithm gives us an accuracy of 0.2 pixels. Each display pixel is about $2.5mm$ in size. Even with this accuracy, however, a misalignment of even a couple of pixels in the reconstructed luminance response can cause perceived discontinuities without the edge attenuation. The edge attenuation alleviates the situation, and we can tolerate greater errors of about 5 – 6 pixels in the display space.

Sampling Density: Sampling density decides the accuracy of the reconstruction of the luminance surface for the display. As is clear from the results, having two times the resolution of the display is ideal and would get rid of any sampling artifacts. More important however, is the minimum sampling density required to reconstruct the surface correctly. Obviously, this will be different from wall to wall. But to get an approximate idea, we did the following experiment. We reconstructed the luminance response of a four-projector region of the wall sampled at the ideal sampling density. The frequency content of the luminance of this region is representative of that of a larger display because the larger display is made of several such four-projector configurations. Fourier analysis of this luminance image after edge attenuation showed that the required sampling resolution is about one-fifth of the display resolution. In our fifteen-projector implementation, the wall is sampled at one-third the display resolution, and still we see some artifacts since there may be places in the wall that were not properly represented by the small region we used to decide on the minimum sampling density.

Dynamic Range of the Calibration Images: It is important for the brightness of each projector to be well within the dynamic range of the camera. This can be verified by simple under or over-saturation

tests of the camera images. In display walls made of many projectors there may be large brightness variation across projectors. In such cases, the camera exposure should be adjusted to accommodate for this variation. This change in exposure should be taken into account by appropriate scaling factors while generating the luminance surface [3]. Using same exposure for all projectors leads to contouring artifacts as seen in the right-most projector in Figure 3 of the color page.

Camera Properties: It is important for the camera not to introduce additional luminance variation beyond that is already present in the wall. Hence, the camera must produce flat fields when it is seeing a flat color. As is mentioned in [3], most cameras satisfy this property at lower aperture settings, especially below $F8$. Our camera had a standard deviation of 2 – 3% for flat field images. These flat field images were generated by taking pictures of nearly diffused planar surfaces illuminated by a studio light with a diffusion filter.

Black Offset: In our method we assume that black produces zero luminance. This is not true in case of the projectors. Because of several leakages in the light path, the projectors have a non-zero black luminance called the *black offset*. Hence, if the image content is near black, we can see faint seams. From our experience, we find that the black offset has less effect on images with high frequency contents.

White Balance: Our method generates a per channel LAM for every pixel. Since each channel may get attenuated differently, the grays may not be retained as grays when transformed by the LAM. Faint color blotches may therefore appear in the results. Hence, we use the LAM generated for the green channel for all channels. Since the nature of luminance variation is similar across the three channels, the small inaccuracy introduced by this does not show any visible artifacts.

4 Conclusion

In this paper we first identify the different parameters that causes intra and inter projector color variation. We study how the changes in these parameters change the luminance and chrominance characteristics of the projectors and hence of the display wall. This gives us the useful insight that for projection based display walls made of same brand projectors luminance variation is the primary cause

of color variation while chrominance variation is negligible. Thus, we infer that luminance matching may be sufficient for achieving photometric seamlessness.

We demonstrate this by presenting a camera based method to achieve photometric uniformity in multi-projector displays by correcting for the luminance variation alone. Our one time calibration procedure generates a luminance attenuation map which is then used to correct any images. The LAM achieves a luminance matching across all the display pixels.

We believe that this is the first step towards achieving photometric uniformity across projection based displays but much more still needs to be done. Following are some of the things we are working on currently.

- Though this method removes the seams, the dynamic range of the display reduces dramatically since we are matching the response of all the pixels to the response of the worst pixel. This leads to under-utilization of system capabilities, especially in the overlap regions which have higher brightness and range. We are currently developing algorithms which can remove seams and at the same time make better use of the resources thus leading to higher dynamic range displays.
- It is important to evaluate the results of algorithms based on some photometric or perceptual metric. This metric should quantify the different display properties that are improved or degraded by the proposed algorithms. We are in process of designing such metrics.
- It is evident that as we move towards bigger display walls, the limited camera resolution will be insufficient to sample the luminance surface adequately leading to sampling artifacts in the corrected images. Hence, there is a necessity to design scalable solutions that can correct parts of the wall at a time and then stitch together the results. We are also investigating such scalable algorithms.
- The proposed method does not depend on the image content. But, if content of the image is considered as an input to the algorithm, the compression in the dynamic range can be reduced leading to higher dynamic range images. We are investigating such content based corrections which may be more suited for canned movies as is used for entertainment purposes.

5 Acknowledgments

We acknowledge Kodak for providing us with high resolution pictures to test our algorithm. This work was supported in part by the U.S. Department of Energy (DOE), under Contract W-31-109-Eng-38.

References

- [1] Ian Buck, Greg Humphreys, and Pat Hanrahan. Tracking graphics state for networked rendering. *Proceedings of Eurographics/SIGGRAPH Workshop on Graphics Hardware*, 2000.
- [2] Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. Defanti. Surround-screen projection-based virtual reality : The design and implementation of the cave. In *Proceedings of ACM Siggraph*, 1993.
- [3] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. *Proceedings of ACM Siggraph*, pages 369–378, 1997.
- [4] D.H. Brainard. Calibration of a computer controlled color monitor. *Color Research and Applications*, 14(1):23–34, Feb 1989.
- [5] D.L. MacAdam. Visual sensitivity to color differences in daylight. *Journal of Optical Society of America*, pages 247–274, 1942.
- [6] Edward J. Giorgianni, Thomas E. Madden. *Digital Color Management: Encoding Solutions*. Addison-Wesley, 1998.
- [7] Mark Hereld, Ivan R. Judson, and Rick Stevens. Dottyfoto: A measurement engine for aligning multi-projector display systems. *Argonne National Laboratory preprint ANL/MCS-P958-0502*, May, 2002.
- [8] Greg Humphreys, Ian Buck, Matthew Eldridge, and Pat Hanrahan. Distributed rendering for scalable displays. *Proceedings of IEEE Supercomputing*, 2000.
- [9] Greg Humphreys, Matthew Eldridge, Ian Buck, Gordon Stoll, Matthew Everett, , and Pat Hanrahan. Wiregl: A scalable graphics system for clusters. *Proceedings of ACM SIGGRAPH*, 2001.
- [10] Greg Humphreys, Mike Houston, Ren Ng, Randall Frank, Sean Ahem, Peter Kirchner, and James Klosowski. Chromium : A stream processing framework for interactive rendering on clusters. *ACM Transactions of Graphics*, 2002.
- [11] G. Humphreys and P. Hanrahan. A distributed graphics system for large tiled displays. In *Proceedings of IEEE Visualization*, 1999.

- [12] Aditi Majumder, Zue He, Herman Towles, and Greg Welch. Achieving color uniformity across multi-projector displays. *Proceedings of IEEE Visualization*, Oct 2000.
- [13] Aditi Majumder. Properties of color variation across multi-projector displays. *To appear in Proceedings of Eurodisplay*, 2002.
- [14] Charles Poynton. *A Technical Introduction to Digital Video*. John Wiley and Sons, 1996.
- [15] Ramesh Raskar. Immersive planar displays using roughly aligned projectors. In *Proceedings of IEEE Virtual Reality 2000*, 1999.
- [16] R. Raskar, M.S. Brown, R. Yang, W. Chen, H. Towles, B. Seales, and H. Fuchs. Multi projector displays using camera based registration. *Proceedings of IEEE Visualization*, 1999.
- [17] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future: A unified approach to image based modeling and spatially immersive display. In *Proceedings of ACM Siggraph*, pages 168–176, 1998.
- [18] Samanta Rudro, Jiannan Zheng, Thomas Funkhouser, Kai Li, and Jaswinder Pal Singh. Load balancing for multi-projector rendering systems. In *SIGGRAPH/Eurographics Workshop on Graphics Hardware*, August 1999.
- [19] Maureen C. Stone. Color balancing experimental projection displays. *9th IS&T/SID Color Imaging Conference*, April, 2001.
- [20] Maureen C. Stone. Color and brightness appearance issues for tiled displays. *IEEE Computer Graphics and Applications*, March, 2001.
- [21] Russell L. De Valois and Karen K. De Valois. *Spatial Vision*. Oxford University Press, 1990.
- [22] Ruigang Yang, David Gotz, Justin Hensley, Herman Towles, and Michael S. Brown. Pixelflex: A reconfigurable multi-projector display system. *Proceedings of IEEE Visualization*, 2001.